

PROCESS CONTROL DIAGNOSTICS

F. Greg Shinskey
Process Control Consultant
North Sandwich, NH 03259

Abstract

With all the tuning methods documented, it is remarkable how often controllers are mistuned, focusing on set point response when the set point is never changed, as in a level loop. Guidelines are provided on recognizing poor load response and correcting it. However, some loops oscillate all the time, or some of the time, and no amount of tuning will correct this behavior. Based on waveform analysis, the nature of the controlled variable and the valve type, the cause can be readily diagnosed and corrected. Liquid-level loops are especially vulnerable to several common maladies including deadband and hydraulic resonance. Special considerations apply to cascade loops; rules are also given for effective filtering.

Introduction

Many control loops fail to provide the tight, stable regulation needed to maximize the performance of the process plants where they are used. There is a variety of reasons for these failures, from ineffective tuning to valve deadband, but all can be diagnosed and corrected, given enough information on behavior patterns and loop configuration. This paper describes many of these problems, their causes and the remedies which have been proven to correct them. Each diagnosis requires information from several sources:

1. What is the problem: poor regulation, set point overshoot, or cycling?
2. If cycling, what is the waveform of the controlled variable: sine, triangular, etc.?
3. What variable is being controlled: flow, pressure, level, temperature, composition?
4. What is the configuration: single-loop or cascade?
5. What is the valve characteristic, and does it have a positioner?
6. What process is being controlled: static mixer, heat exchanger, exothermic reactor?

Based on the author's experience and the answers to the above questions, most loops can be optimized.

Poor Load Regulation

A common reason for poor load regulation is the practice of tuning controllers for set point changes, which is taught in the universities. Some academics even complain that the tuning rules developed by Ziegler and Nichols¹ produce excessive set point overshoot. A careful reading of their work, however, reveals no set point changes made. They simulated step load changes by biasing the controller output, and their rules optimize settings for step load response. Load regulation is far more important than set point response, considering that in continuous processes, set points on most critical loops such as liquid level, pressure, temperature, and pH, are rarely changed, while load changes are frequent and can be severe.

Some controllers have an adjustable output bias which can be stepped to simulate a load change. Otherwise, one can be simulated by transferring the controller to manual while in the steady state with no deviation, by stepping the controller output, and immediately transferring back to automatic before a deviation begins.

To identify the characteristics of optimum load regulation, a distributed lag has been simulated, representing processes such as heat exchangers, distillation columns, and stirred tanks. The curve in Fig. 1 describes its open-loop step response. The 63.2-percent response point occurs when time t equals the total time response τ of the system. The optimum settings for a Proportional-Integral (PI) controller on this process are:

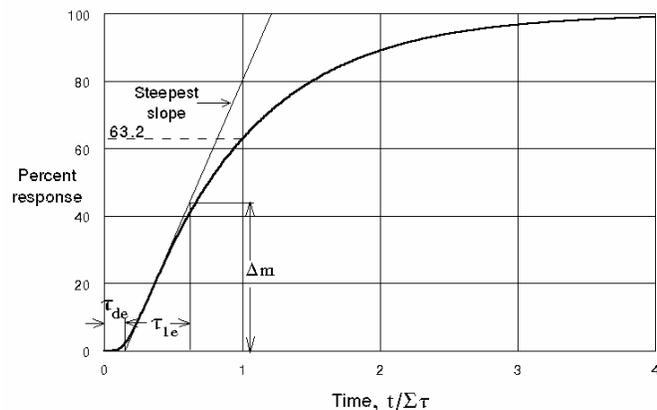


Fig.1. Identifying a distributed-lag process through its step response.

$$P_{\text{opt}} = 18 K_p \quad I_{\text{opt}} = 0.58 \Sigma \tau \quad (1)$$

where P_{opt} is the optimum proportional band in percent, I_{opt} is the optimum integral time expressed in the same units as τ , and K_p is the steady-state process gain.

The principal difficulty is in identifying K_p and τ for slow processes like distillation columns. To determine them directly from the curve, about 99-percent response must be reached, which requires an elapsed time of 4 τ in the absence of any disturbances. For a distillation column of 50 trays, 4 τ could be as long as 7 hours. Ziegler and Nichols used a faster method, based on the slope of the curve at its steepest, and the intercept of that slope, as shown in Fig. 1. The estimated deadtime, τ_{de} , is the time intercept. The slope is estimated in terms of the time τ_{1e} required for the controlled variable to change an amount equal to the size of the step input m , when moving at its fastest rate. (In the example used for Fig. 1, K_p happens to be >1 .)

Because distributed lags all have the same shape curve, the estimated properties can be directly converted into the defining parameters:

$$K_p = 7.5\tau_{de} / \tau_{1e} \quad \Sigma \tau = 7.0\tau_{de} \quad (2)$$

The dashed curve in Fig. 2a is a reproduction of the open-loop step response, which is also the step load response under manual control. The center of the three solid curves is the optimum load response for a PI controller for the same change in load. The optimum curve has the minimum Integrated Absolute Error (IAE) of all possible combinations of P and I settings, a criterion which represents the best combination of minimum peak deviation, minimum integrated error, and effective damping. The minimum-IAE curve may be recognized as having a symmetrical first peak similar to a Gaussian-distribution curve, a slight overshoot, and low decay ratio.

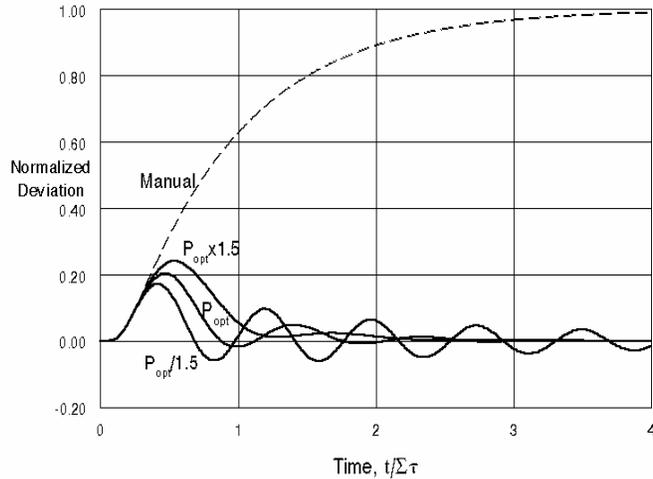


Fig. 2a. Step load responses with different proportional bands.

To help in diagnosing ineffective tuning, Fig. 2a also describes the step load response with two suboptimal proportional-band settings. When the proportional band set at 1.5 times the optimal, the response curve loses its symmetry, requiring longer to return to set point following the first peak than to reach the first peak. This unsymmetrical response is very common in control loops where operators emphasize stability over recovery. If loop gain is variable, this behavior can be observed under conditions of low gain.

Figure 2a also shows the results of setting the proportional band at the optimum divided by 1.5: recovery is so fast that the set point is overshoot, and the damping is very light. Note also that the peak height varies directly with the proportional band.

The effects of suboptimal integral settings can be distinguished from those of suboptimal proportional settings. Figure 2b shows that integral action primarily affects the overshoot of set point, without relocating the first peak. Setting the integral time at 1.5 times the optimal actually causes undershoot, and at the optimal divided by 1.5 exaggerates the overshoot. To determine whether cycling is due to an improper integral setting, consider that for this loop, the ratio of the observed period to the

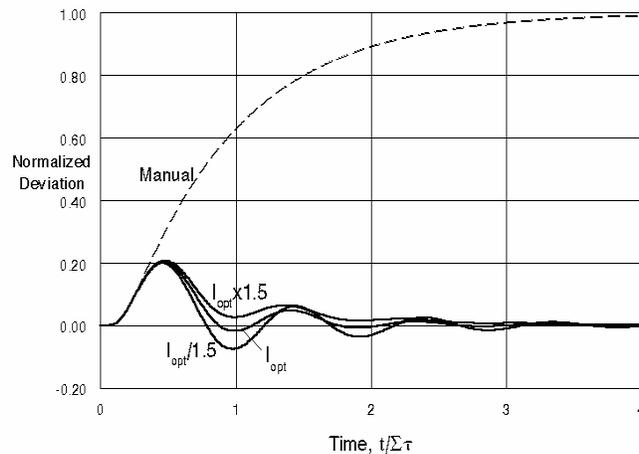


Fig. 2b. Step load responses with different integral times.

optimum integral time is 1.62. The curve with the undershoot has a ratio of 1.02, and the curve with the excessive overshoot has a ratio of 2.53.

Derivative action is underutilized by industry, probably because it adds a third dimension to the tuning problem. To add to the confusion, there are two fundamentally different configurations of industrial PID controllers—interacting and noninteracting—and their tuning rules are different. To reduce the dimensionality of the tuning procedure, it is suggested that the integral and derivative time settings should be adjusted together in a fixed ratio. The optimum settings for the noninteracting PID controller are:

$$P_{\text{opt}} = 10K_p \quad I_{\text{opt}} = 0.30\sum\tau \quad D_{\text{opt}} = 0.09\sum\tau \quad (3)$$

and for the interacting version are:

$$P_{\text{opt}} = 15K_p \quad I_{\text{opt}} = 0.25\sum\tau \quad D_{\text{opt}} = 0.10\sum\tau \quad (4)$$

The proper use of derivative action results in a period that is approximately half that of the PI controller, and a corresponding reduction in the integral setting as well. For the distributed lag, integrated error per unit load change, which varies directly with the product of proportional band and integral time, is reduced from that of a PI controller by a factor of 2.7 by the interacting PID and 3.4 by the noninteracting version. The optimum ratio of period to integral time for the interacting PID is 2.4, and for the noninteracting PID is 1.62, which is the same as for the PI controller.

Set point Overshoot

If the PID settings are optimized for load response, then a step change in set point will produce a large overshoot on lag-dominant processes. However, if they are adjusted to minimize set point overshoot, achieved with a longer integral time and wider proportional band, the result will be a slow exponential recovery from load changes. Some controllers feature the option of eliminating proportional action on the set point (none should have derivative action on the set point), which produces a set point undershoot, sometimes desirable. However, some modern controllers have instead, a lead-lag set point filter. Typically, the lag setting of this filter is linked to the integral setting, and the only adjustable term is the lead/lag ratio. This results in applying a proportional gain to the set point which is a fraction of the proportional gain of the controller, the fraction adjustable over the range of 0-1. After the PID settings have been optimized for load response, the lead/lag ratio—the dynamic gain of the filter—is adjusted to optimize set point response. Settings in the range of 0.25 to 0.5 are used for lag-dominant processes. Compare this to an effective value of zero for the controller with no proportional action on the set point.

Filtering is only effective when the controller does not saturate. Large set point changes causing saturation, or startup of a process with a large initial deviation, may result in overshoot which a filter cannot correct. The cause of the overshoot in these situations is integral windup, and the correction is applied to the integral term. It usually takes the form of logic which disables integral action when an output limit is reached, replacing the constant of integration with an adjustable “preload” setting. Special consideration applies to the primary controller in a cascade system, which will windup when the output of the secondary controller reaches a limit. The

preferred solution to this problem is to connect the secondary controlled variable to the integral feedback loop of the primary controller, as described later.

Continuous Cycling

Some loops cycle all the time, and adjustment of the mode settings only affects the period and amplitude of the cycle. A constant-amplitude cycle is called a “limit cycle,” and is due to a nonlinear element in the loop. The most common limit cycle appears in liquid-level and some gas-pressure loops, caused by deadband in the control valve. Deadband is the difference between the stem position of a valve and its control signal as a function of whether that signal is increasing or decreasing. If, for example, deadband causes the stem position to be 3 percent below the control signal by while the signal is increasing, it will also cause it to be 3 percent above on a decreasing signal, and to stop moving altogether when the difference is less.

In a level loop with deadband and a PI controller, the valve position and hence the flow will cycle in a clipped sine wave, while the level cycle is almost triangular, as shown in Fig. 3. The period of the level cycle in the units of integral time is:

$$\tau_o = I(5 + \sqrt{P}) \quad (5)$$

and its amplitude in percent is:

$$A = 1.2aP/100 \quad (6)$$

where a is the deadband in percent. These relationships are not analytical, but were fit to simulation results. The simulation in Fig. 3 has a 5-percent deadband in the valve, appearing as the difference between the peak values of controller output and stem position. Doubling the proportional band of the PI controller at time 240 s increased both the amplitude and the period of the cycle, as described by the above equations.

A digital blending system wherein accumulated volume is the controlled variable, is subject to the same limit cycling as a liquid-level loop. A gas-pressure loop differs from liquid level in being self-regulating where level loops are not. While it can limit cycle in a triangular wave, it responds somewhat differently to the proportional setting. In a gas-pressure loop, increasing the proportional band increases the period but *reduces* the amplitude of the oscillation. Unfortunately, however, this also increases the peak deviation following a load change.

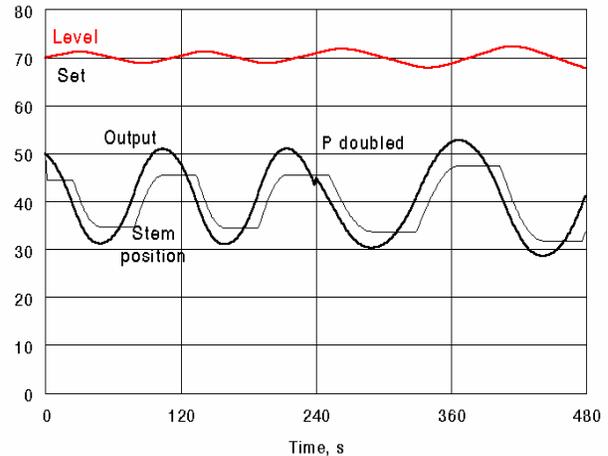


Fig. 3. A level loop limit-cycling; P was doubled at 240 seconds.

Deadband is caused by friction in the valve seals and guides. It can be overcome by using a valve positioner to close the loop around the valve actuator. However, if a *volume booster* is inserted between the positioner and the actuator, the deadband may return. For critical gas-pressure loops, as on a compressor, the use of a *smart digital positioner* is recommended. For critical level loops, cascading to a secondary flow controller is even better, as this both closes the loop around the valve and linearizes its behavior, as well as canceling the effects of supply-pressure changes and interactions with other flow loops. But *do not* cascade gas pressure to flow on a compressor—their periods are too close together for dynamic stability.

Flow and liquid-pressure loops will *not* limit-cycle in the presence of deadband, and therefore do not require a positioner on their valve. In fact, the use of a valve positioner in these loops is *destabilizing*. Because the valve actuator is the slowest element in both the position and the control loop, their periods are almost identical. Stabilizing the control loop requires increasing the proportional band of the controller excessively, which degrades the load response of the loop by the same factor.

Averaging Level Control

There are some level loops in the plant that are intended to buffer changes in flow rather than to control level tightly. Tight level control passes swings in flow entering the tank directly to the downstream process, and this is undesirable where the downstream process is sensitive to variations in feed rate. However, simply loosening level control by widening the proportional band while retaining integral action is not recommended, as it actually *destabilizes* the loop.

A PI controller closing the loop around an integrating (non-selfregulating) process forms a second-order underdamped system having a period of

$$\tau_n = 2\pi(\tau_i IP/100)^{0.5} \quad (7)$$

where τ_i is the time constant of the integrating process. The damping factor is similarly related:

$$\zeta = 0.5(100I/P\tau_i)^{0.5} \quad (8)$$

Observe that increasing the proportional band *reduces* the damping factor while increasing the period of oscillation. This is surely counter-intuitive, and has led in many plants to cyclical swings in feed rate having periods of an hour or more.

The solution to this problem is to avoid using integral action. There is no advantage to returning the level of a surge tank to a fixed (set) point in the steady-state², so integral action is not required, and in fact is a hindrance. Level should be allowed to change with load, as this maximizes the buffering action of the vessel. The band of the proportional controller should be set to the maximum allowable level variation, typically 70-90 percent. Proportional control,

however, creates a problem of perception in the minds of the operators: the appearance of set point offset leads to a conclusion that the controller is malfunctioning. A way around this problem is to use a calculation block instead of a controller, programmed as follows:

$$m = (c - c_l)/(c_h - c_l) \quad (9)$$

where m is the signal to the valve, c is the measured liquid level, and subscripts l and h identify its desired low and high limits respectively. A first-order filter as large as $0.4 \tau_i$ can be inserted in the signal to the valve for additional smoothing without destabilizing the loop.

Flow-dependant Stability

Some control loops change their behavior as a function of flow. There are several possible reasons for this, the most common being the use of the wrong valve characteristic. Equal-percentage valves are over-prescribed, often used on loops where a linear valve would be a better fit. The equal-percentage characteristic is logarithmic, its gain increasing directly with flow passed, as shown in Fig. 4 for a pressure-drop (dp) ratio of 1.0. If a loop tends to oscillate at high flow and is sluggish at low flow, it probably has an equal-percentage valve and shouldn't.

The variable gain of an equal-percentage characteristic is often used to offset the effects of variable pressure drop resulting from pump head drop and line resistance, but it overcorrects. Figure 4 shows that when the pressure drop at full flow is a factor of 0.33 of that at no flow (dp ratio of 0.33) the installed characteristic of an equal-percentage valve is about as far out of linear as that of a linear valve. Unfortunately, there is little other choice in valve trim. Smart digital positioners allow the selection of any desired characteristic, and can correct for this.

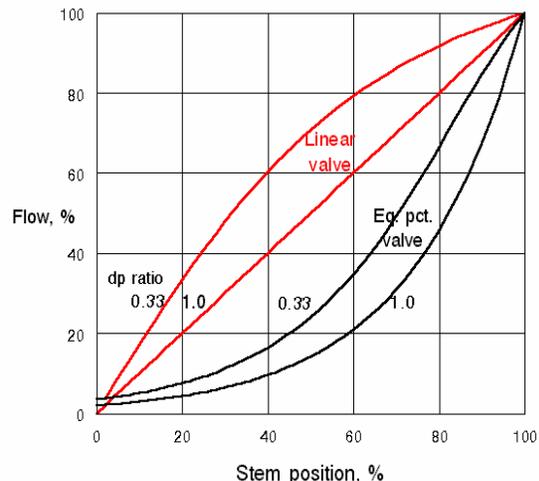


Fig. 4. At a dp ratio of 0.33, both valves are equally nonlinear.

Flow, level, and pressure loops should be fitted with linear valves. Temperature loops in general control better with equal-percentage valves, in that most have a steady-state gain which varies inversely with flow. *Heat exchangers* are particularly problematic in this regard, in that their τ_i also varies inversely with flow. Therefore, if PID settings are to be optimum at all flows, all three must vary inversely with feed rate, when manipulating a flow loop in cascade. Reference 3 describes the application of a PID-deadtime controller to a superheater in a power plant, where all four controller settings were programmed as functions of steam flow.

Static mixers behave in the same way as heat exchangers, for the same reason—no recirculation, as there is in a stirred tank. A reduction in feed rate of only 35 percent of value can bring the PI composition loop from optimum response to instability—this is its measure of *robustness*. An equal-percentage valve can help, but only provides gain correction—dynamics also vary inversely with flow. The best fix for a static mixer is to put a *recirculation pump* around it—this minimizes its deadtime, allowing smaller values of both proportional band and integral time, which are then optimum for all feed rates.

Limited valve *rangeability* can cause limit-cycling at very low flow rates. If the process load ever demands less flow than the lower controllable limit of the valve, the valve will tend to close completely. But since zero flow is too little to satisfy the load, the controller will subsequently open the valve, delivering too much flow and causing the controller to close it again. The result is a sawtooth cycle such as that shown in Fig. 5. If the load happens to require exactly half the minimum controllable flow, the cycle may be a symmetrical triangle, but otherwise it will be unsymmetrical or sawtooth. This limit cycle is more likely to develop if the valve is oversized—the remedy is to resize the valve if that is the problem. If it is not oversized, a second, smaller valve may be sequenced with the first to extend overall rangeability; this is best achieved using equal-percentage valves, and a smooth characteristic requires that they *not* both be open at the same time⁴.

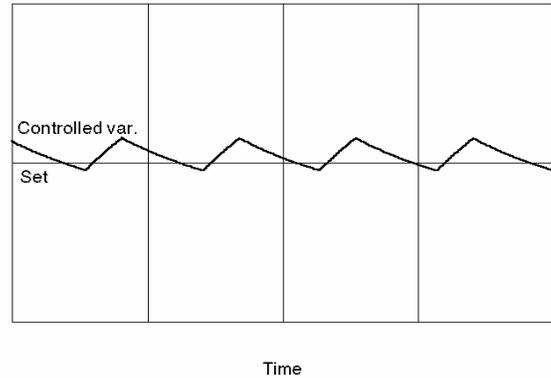


Fig. 5. A sawtooth wave caused by load below the throttling limit of the valve.

Amplitude-dependant Stability

This can be a particularly severe problem, in that a loop that is normally stable at set point can explode into expanding oscillations following a large upset, until the output of the controller is cycling between its limits. Return to stability requires manual intervention. Operators develop a fear that it will happen again, and the conditions which caused it may not be precisely known.

The usual cause of this behavior is a controller integrating faster than its control valve can follow. Valves are all velocity-limited, able to follow a small change much faster than a large one. For example, a valve that requires ten seconds for full stroke, can follow a ten-percent signal

change in only one second, and its controller is likely to be tuned based on that small-signal response.

Figure 6 describes a pressure-control loop on a compressor, where the integral time of the controller is set at half the stroking time of the valve. Response to a 25-percent load change is stable, but a 26-percent step and larger cause an expanding cycle. In this simulation, the stroking time of the valve is 10 s and the integral time of the controller is 5 s. If the integral time is set equal to the stroking time of the valve, the loop is stable to all upsets, although several cycles will develop following larger upsets, and integrated error will double for all. The preferred solution to the problem is to speed up the valve: use a smart positioner—not a volume booster. If the valve speed cannot be increased, then a velocity limiter matched to the valve should be placed in the integral feedback loop of the controller, so that it cannot integrate faster than the valve can stroke.

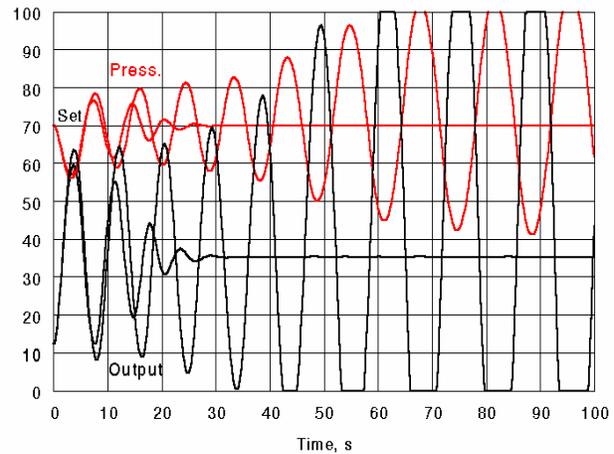


Fig. 6. This loop recovers from a load change of 25%, but not 26%.

pH Loops

Control of pH poses special problems, because of the nonlinear and very sensitive relationship between the controlled variable and the flow of reagent. If the effluent pH measurement is very noisy, the reagent is not being effectively mixed and distributed with the feed—a difficult problem because their flow ratio can be as great as one part in 20,000. Reagent should be premixed with the feed at its entrance, and the entire contents of the neutralization vessel circulated by a high-velocity agitator, with vortex formation avoided by proper baffling. The effluent pH measurement should be made at the exit of the vessel, but within the mixed zone, so that it will be representative of the effluent, and responsive to the reagent flow, even when the feed rate is zero.

Valve positioners should always be used in pH control, to avoid deadband. In addition, the reagent lines should terminate in a *loop seal*, so that flow starts as soon as the valve opens and stops as soon as it closes, without the transfer line having to fill or drain.

When controlling pH in the neutral range, which is almost always the case, a curve characterizer is recommended at the input to the controller. It can be applied separately to measurement and set point, or singly to the deviation. A simple error-square function is often sufficient—it doesn't have to be a perfect fit. For industrial wastewaters, the titration curve is often variable, changing the loop gain widely. Limit-cycling can develop when buffers are absent, and heavy damping ensue when they return. The only remedy for this problem seems to be a self-tuning controller, but nonlinear characterization is still recommended to keep the controller from retuning on every load change.

The user should be aware that pH sensors have some strange failure modes. For example, an indicated pH of 7 is not necessarily good, because a shorted electrode produces the same result. Contamination of the reference electrode usually causes a bias, as does a ground current rectified by the electrode. A slowly fouling sensor increases its time constant as the film thickness builds. This causes the period of the loop to increase gradually over time, and the amplitude as well. If this happens infrequently, chemical cleaning is recommended; if frequently, ultrasonic or mechanical cleaning may have to be installed to remove the film periodically or prevent its growth.

Cascade Control

Cascade control is recommended for loops with complex dynamics such as composition control of distillation, and temperature control of exothermic reactors. The secondary loop can reject disturbances originating within it, often before they have any effect on the primary variable. Closing the loop around secondary dynamic elements also speeds up the primary loop. In addition, any nonlinearities closed within the secondary loop are thereby removed from the primary loop. As an example of the latter, consider controlling a reactor temperature by manipulating the flow of cooling water. Heat transfer into the cooling system varies nonlinearly with flow, and also with both reactor and coolant temperatures. When the reactor temperature controller sets coolant exit temperature in cascade, however, the primary loop is linear: every degree change in coolant temperature produces the same change in reactor temperature.

Because the output of the primary controller sets the set point of the secondary controller, some engineers may then assume that the secondary controller should be tuned to optimize its set point response, either through its PID settings, a set point filter, or by elimination of proportional action from its set point. This is a big mistake! Optimum performance for the primary loop depends on the secondary controller being tightly tuned for load rejection, with *no filter* between the controllers. In fact, using a secondary controller without proportional action on the set point is equivalent placing a first-order lag equal to the secondary integral time between the two controllers. The performance of the primary loop then becomes *worse* than if it were operating as a single loop—cascading is then a liability rather than an advantage.

In the discussion of integral windup above, it was pointed out how the primary controller in a cascade system could be protected by connecting the secondary controlled variable as primary integral feedback. This is shown schematically in Fig. 7 for an interacting PID primary controller, where integration is achieved by positive feedback of the output through

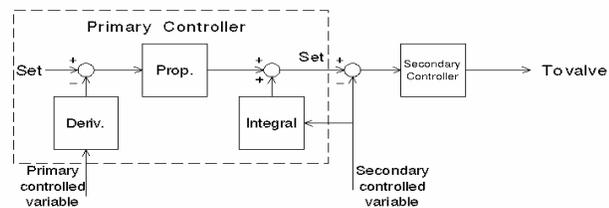


Fig. 7. Configuring a cascade system for maximum effectiveness.

a first-order lag representing the integral time. This configuration allows the primary controller to be left in automatic all the time, even when the secondary is in the manual mode, thereby simplifying startup and manual operations. Furthermore, it places the *entire secondary loop* within the integral loop of the primary controller. This has two benefits: the primary integral time can be reduced somewhat and thereby give tighter control, and variations in the dynamics of the secondary loop also vary the rate of integration of the primary controller, increasing robustness substantially. During pilot-plant tests of this configuration, temperature control of an exothermic batch reactor which set coolant exit temperature in cascade required no retuning while the primary time constant of the reactor was varied over a 4:1 range.

By contrast, the temperature of an exothermic reactor under conventional cascade control was found to limit cycle while the heating valve was locked out, because the primary controller wound up every time the secondary controller tried to open the locked-out valve. Later, when the configuration was changed to use secondary temperature as the primary integral feedback signal, the limit cycle disappeared.

Noise and Resonance

Noise is a random fluctuation in the measured variable having a frequency content above the bandwidth of the control loop. As such, it cannot be controlled by the loop, and when acted upon by the proportional and derivative modes of the controller, only serves to overdrive the valve and even amplify itself. Derivative action cannot be used in its presence, which is why derivative is never used in flow and level control. The proportional band of the controller may be widened to save the valve, but at the sacrifice of load response. The alternative is to insert a filter on the measured variable.

It must be recognized that filtering is counterproductive to control: a filter reduces the motion of the controller output in response to input variations, whereas the controller attempts to move its output to reduce input variations. In this light, the optimum amount of filtering is the least that is acceptable. In the act of attenuation, the filter also introduces phase lag, which requires that the controller be detuned in its presence—process deadtime is essentially augmented by the filter time. It also increases the integrated error IE following a disturbance:

$$IE = \Delta m(P/100)(I + \tau_f + \Delta t) \quad (10)$$

where m is the change in controller output between any two steady states, τ_f is the time constant of the filter, and t is the sample interval of the (digital) controller. (For the PID-deadtime controller mentioned earlier, the deadtime setting replaces t in the above expression.)

Some liquid-level loops exhibit hydraulic *resonance* as waves oscillating between bounded surfaces such as the sides of a tank or the opposite ends of a U-tube. The resonant period is a function of the distance between the bounded surfaces, following the same formula as that of a pendulum. For example, 30-ft-long feedwater heaters in a power plant produced a 6-s resonant cycle in the level of condensate within, so lightly damped that very little proportional action by the level controller caused sustained oscillations. Attenuating these cycles by a factor of 10 with

a filter allowed the proportional band of the level controller to be reduced by a factor of 5, greatly stabilizing control.

A first-order filter may be used to suppress either noise or resonance, but a second-order Butterworth filter is more effective. Given the need to attenuate a resonance by a factor of 5: at a period 3 times that of the resonance, a second-order Butterworth filter will pass 89 percent of the signal, whereas a first-order filter will pass only 54 percent of the signal. In other words, the Butterworth has a sharper cut-off.

An ideal filter for resonance is the moving-average filter. It averages the last n samples arithmetically, where filter time $n.t$ is selected to match the period of the resonance. If the match is perfect, the resonance will be completely eliminated. In response to a step change, this filter produces a ramp exactly $n.t$ long.

To keep a filter from severely interfering with the action of the controller, it should be set $<D/10$ for a PID controller, and $<I/5$ for a PI controller.

Sampling

Sampling is common to all digital controllers, and to some measurements such as chromatographic analyzers. Information theory warns about aliasing errors introduced when the sample interval t approaches the period of any cycle in the signal. But in a closed loop, this presents no danger, because sampling introduces enough phase lag to keep the period of oscillation greater than $2.t$. The problem is instead that sampling augments any deadtime in the loop, just as a filter does, thereby extending the period. While a simple sample-and-(zero-order) hold function introduces only half as much phase lag as an equal deadtime, sampled signals are nearly always filtered over at least one sample interval, which doubles the phase lag. Therefore, when controlling a lag-dominant process, highest performance will be achieved using the shortest sample interval available. This, of course, must be balanced against the demands of other loops and functions running on the same digital processor.

The effective gain of the derivative control mode is limited to D/t . When the derivative gain falls below 10, the mode starts to lose its impact on the loop. Therefore, for the same reasons cited above for filtering, t should be $<D/10$ for PID controllers and $<I/5$ for PI controllers.

Chromatographic analyzers constitute a special case, because of their relatively long sample interval and absence of filtering. The stepwise nature of the controlled variable creates a problem with derivative action in the loop, as it would produce an output pulse proportional to 10 times the input step size at each sample. To obtain most of the effectiveness of derivative action without these pulses, the controller may be synchronized to the analyzer, executing its algorithm only once as soon as a new analysis is available. This lowers the effective derivative gain to the aforementioned D/t , but still obtains some benefit from it. It also provides protection against analyzer failure and variability in the sample interval, because the controller holds its output constant between samples. Alternatively, a filter of $0.4.t$ may be applied to the sampled signal. The performance of these two alternatives is about the same, giving an integrated error about one-third higher than possible under normal PID control, but half that achieved under PI control.

References

1. Ziegler, J. G., and N. B. Nichols, "Optimum Settings for Automatic Controllers," *Trans. ASME*, November 1942, pp. 759-768.
2. Taylor, A.J., and T.G. la Grange, "Optimize Surge Vessel Control," *Hydrocarb. Proc.*, May 2002, pp. 49-52.
3. Shinskey, F. G., "PID-deadtime Control of Distributed Processes," *Contr. Eng. Prac.*, 9 (2001), pp. 1177-1183.
4. Shinskey, F. G., *Process Control Systems*, 4th ed., McGraw-Hill, New York, 1996, pp. 62-64.